ROMPOS · ROMANIAN POSITION DETERMINATION SYSTEM

AI-assisted automation of Bernese GNSS Software processing

PRESENTED BY

Miluță DULEA-FLUERAȘ

The 12th EUPOS® Technical Meeting · Warsaw, October 8th 2026

// 02 THE STARTING POINT

Where we started

One day at a time

Bernese 5.4 driven through its menu on one machine. Every campaign day was started and checked by hand.

Scattered data

RINEX 2 and 3, 30 s and 1 Hz, in 645 folder layouts across archives nobody had fully indexed.

Knowledge in mailboxes

The answers lived in 18 years of IGS Mail, BSW mail and EPN guidelines — impossible to search.

Metadata that drifts

Antenna swaps, relocated monuments and clashing marker names silently break a whole day.

THE QUESTION

What if we could ask the processing system instead of instead of operating it?

...and reprocess every year we have, in one consistent reference frame, while doing it.

// 03 WHERE WE ARE NOW

The result, in numbers

6,015

daily network solutions

2009 – 2026, SINEX + normal equations

756

GPS-week solutions

rigorous ADDNEQ2 combinations, screened

17

years, one reference frame

2009–2026 (no data for 2011)

69

days of processing

1 Jul → 8 Sep; fixes and delivery by 14 Sep

16.5 h

to reprocess a full year

2009, 2010, 2012 — 7 days in parallel

167 GB

delivered, 13,598 files

every file SHA-256 verified

// 04 TIMELINE

Three months, June → September 2026

JUNE

Mid-Jun
Bernese 5.4 in podman containers on odin — 7 days in parallel

30 Jun
AI stack built: state database, MCP tools, agent, RAG, audit log

JULY

2 Jul
Local LLM replaced by the Claude API + live web search

3 Jul
2025 complete — 365 / 365 days

4–12 Jul
2024 & 2023; 69,479 station-days recovered from 1 Hz data

12 Jul →
2016–2022 with CODE repro3 products

AUGUST

10 Aug
Streamlit console replaces the old GUI

19 Aug
Final GPSEST split into clusters: 2.5× faster days

19–27 Aug
Back-processing 2015 → 2012

SEPTEMBER

1 Sep
Whole archive indexed: 15.3 M records in 2 h 14 min

1–8 Sep
2010 and 2009 — the oldest years

5 Sep
Nightly 3-2-1 backup, self-tested

13–15 Sep
QC agents, weekly ADDNEQ2, 2026 fixes, delivery

// PART 01

Breaking the Bernese bottleneck bottleneck

Measure, split, containerise, parallelise

■ AI layers ■ interfaces ■ infrastructure

One stack, eight layers

Interfaces	Streamlit console (9 pages) · natural-language chat · Claude Code for engineering
Guardrails	dry-run by default · human approval for every write · JSONL audit log
Knowledge	RAG: 41 documents, 7,248 chunks (BM25) · sources re-fetched weekly · live web search
Brain	Claude agent — plans, calls tools, explains results in plain language
Hands	MCP server, 8 tools: status · query · quality · knowledge · audit · run a day · decimate · refresh
State	SQLite state index · archive index: 15.3 M records · campaign log
Compute	odin: EPYC 7313, 32 threads, 503 GB · rootless podman · Rocky Linux 8 + Bernese 5.4
Data	RINEX 2 archive · 1 Hz RINEX 3 repository · 30 s store · 14 EPN/IGS ties · CODE products

Where one day’s processing time goes

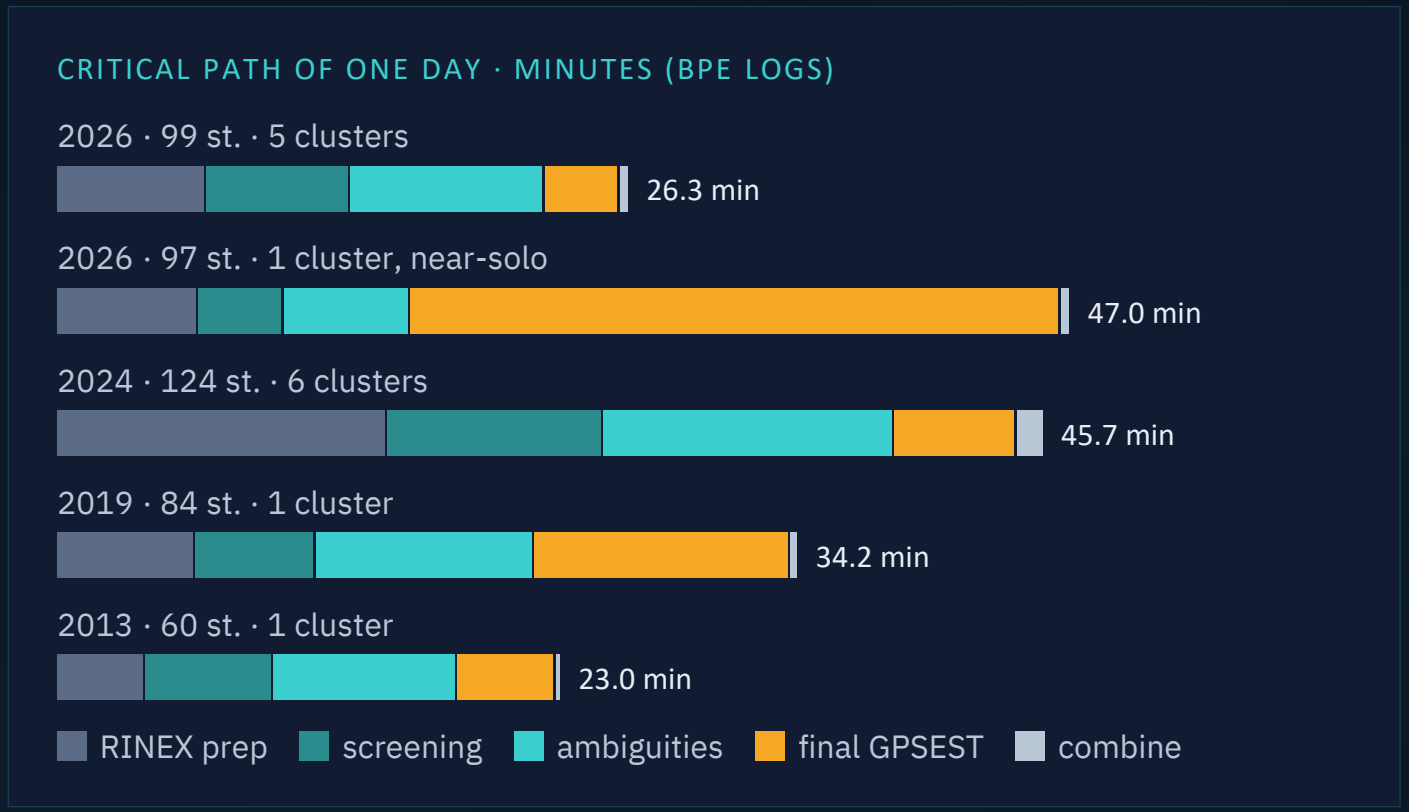

64 jobs

11 parallel — yet one serial chain

the critical path is 99.4–99.9 % of the day’s wall time

1 thread

per Fortran program

GPSEST, RNXGRA, ADDNEQ2: gfortran -O3, no OpenMP — CPU time = wall time

15–39 %

of 32 threads used by one day

≤ 1 job is running for 34–80 % of the day

// 08 THE GPSEST WALL

One program sets the pace — so we split it

MEDIAN DAY DURATION (MIN) BY NETWORK SIZE

FINAL GPSEST · MEDIANS

2026 replay, 7 days in parallel: 28 one-cluster vs 21 five-cluster days

39.7 → 3.4 min

one big GPSEST → several small clusters

70.2 → 28.1 min

the whole day

Why: GPSEST cost grows ≈ baselines^1.5 — small clusters are far cheaper. **How:** V_CLUFIN = A 80 (split above 79 baselines, not 99).

// 09 PARALLEL BERNESE

Rocky Linux containers, one day each

odin · Ubuntu 22.04 · AMD EPYC 7313, 16 cores / 32 threads · 503 GB RAM

rootless podman 3.4.4 · crun · cgroup v2 · overlay storage on NVMe

day 1	day 2	day 3	day 4	day 5	day 6	day 7

each: Rocky Linux 8.9 · Bernese 5.4 · own campaign + private GPSUSER

DATAPOOL 53 GB
shared · staged serially

GLOBAL 93 MB
models · read-only

SAVEDISK → NAS
results · byte-verified

ISOLATION RULES

1. **One container per day**
   its own CAMPAIGN54 and a private 24 MB GPSUSER copy — shared copies race
2. **Stage serially, run in parallel**
   days of one GPS week share orbit products in DATAPOOL
3. **Models read-only**
   GLOBAL mounted :ro; no chown of shared mounts
4. **~2 s per day**
   container start/stop overhead

// 10 THE IMAGE

Rebuilding Bernese as a container image

bernese54-rompos:latest	7.02 GB · 37 layers	
ROMPOS layer Qt4 menu · perl-core · fonts · licence · Hatanaka · autobernese		+0.77 GB
bernese54-headless:rocky8 colleague's base	6.25 GB · 22 layers	
Bernese 5.4 compiled 88 Fortran programs, gfortran 8.5 -O3		4.3 GB
distribution archives + toolchain	1.44 GB + 285 MB	
Rocky Linux 8.9 base image		204 MB

WHAT THE HEADLESS BASE LACKED

the Qt menu program
BPE calls it for every job — Qt4, so cross-built on CentOS 7

perl-core and Hatanaka CRX2RNX
BPE scripts need Time::HiRes; RNX_COP needs CRX2RNX

licence, fonts, GLOBAL models
USER.TXT, fontconfig; DATUM, SATELLIT, ephemerides

a working autobernese CLI
two small patches (terminal size, click-aliases)

Rebuild recipe + 5 GB image tar in the nightly backup

// 11 SCHEDULING & CPU

Keeping 32 threads busy

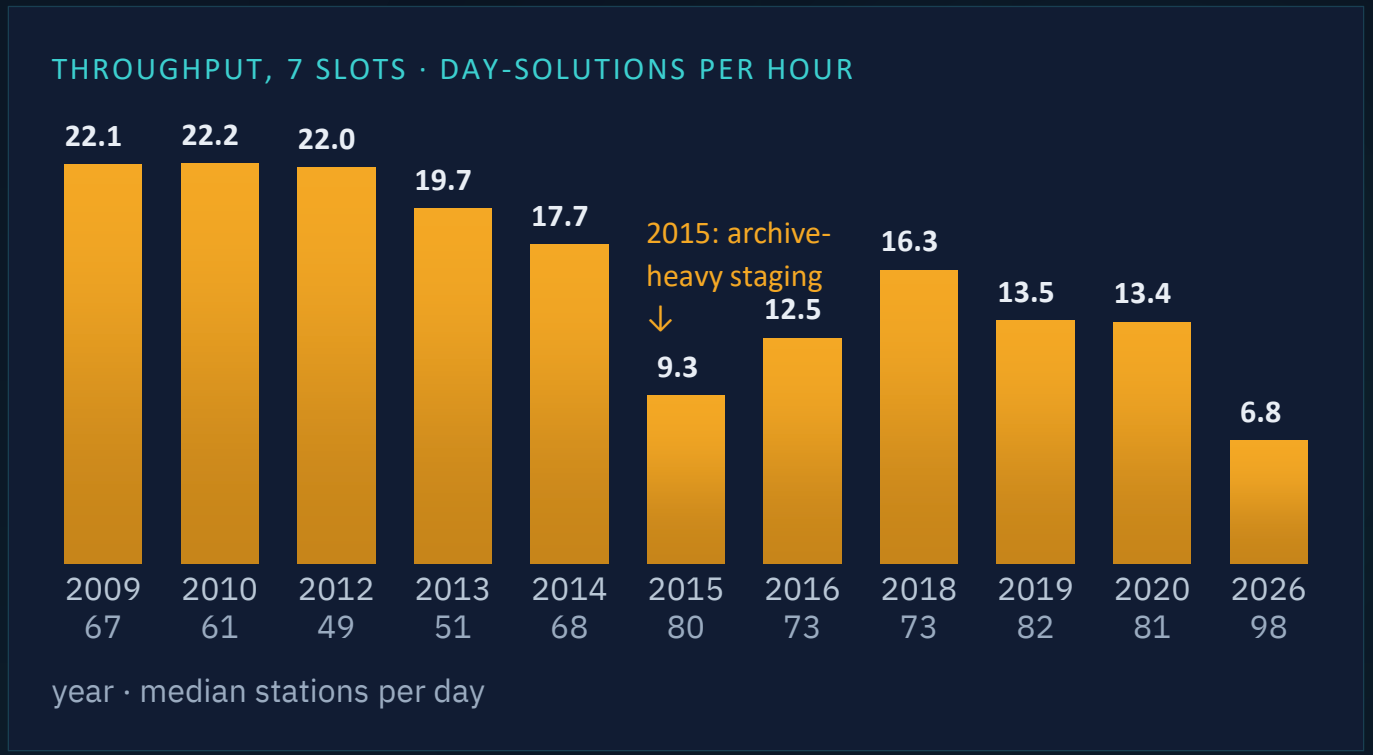

7 × 32

day slots × BPE jobs per day

up to 224 Bernese jobs on 32 threads (68 at once on average, 2013)

2.7–4.5×

throughput vs one day at a time

each day runs 1.7–2.2× slower under load — net win

+14 %

rolling slots vs 7-day batches

no barrier: a free slot takes the next day (2024: 7.7 → 8.8 d/h)

From raw archive to delivery — automated end to end

01
Index
15.3 M records
year from filename
truncated → salvaged

02
Stage
RINEX 2 → 3
marker & antenna fix
+ 14 IGS/EPN ties

03
Process
RNX2SNX BPE
1 container / day
7 in parallel

04
Verify
Helmert vs IGS20
final vs float QC
antenna check

05
Harvest
copy to NAS
byte-verified
local cleanup

06
Combine
GPS-week ADDNEQ2
bad days screened
756 weeks

07
Deliver
one button
timestamped
SHA-256 per file

$ roll_campaign.py 2013 1 365 7

one command runs a whole year — every step logged; nothing is overwritten until the new result is verified.

// PART 02

17 years in 69 days

The campaign, the network and the console

// 14 OPERATIONS SPEED

17 years processed in 69 days

30 min

median network day

6,012 timed days · p10 17 min · p90 52 min

57 min

7 full-network days, in parallel

a week of data in under an hour

16.5 h

per full year

2009, 2010, 2012 (49–67 stations)

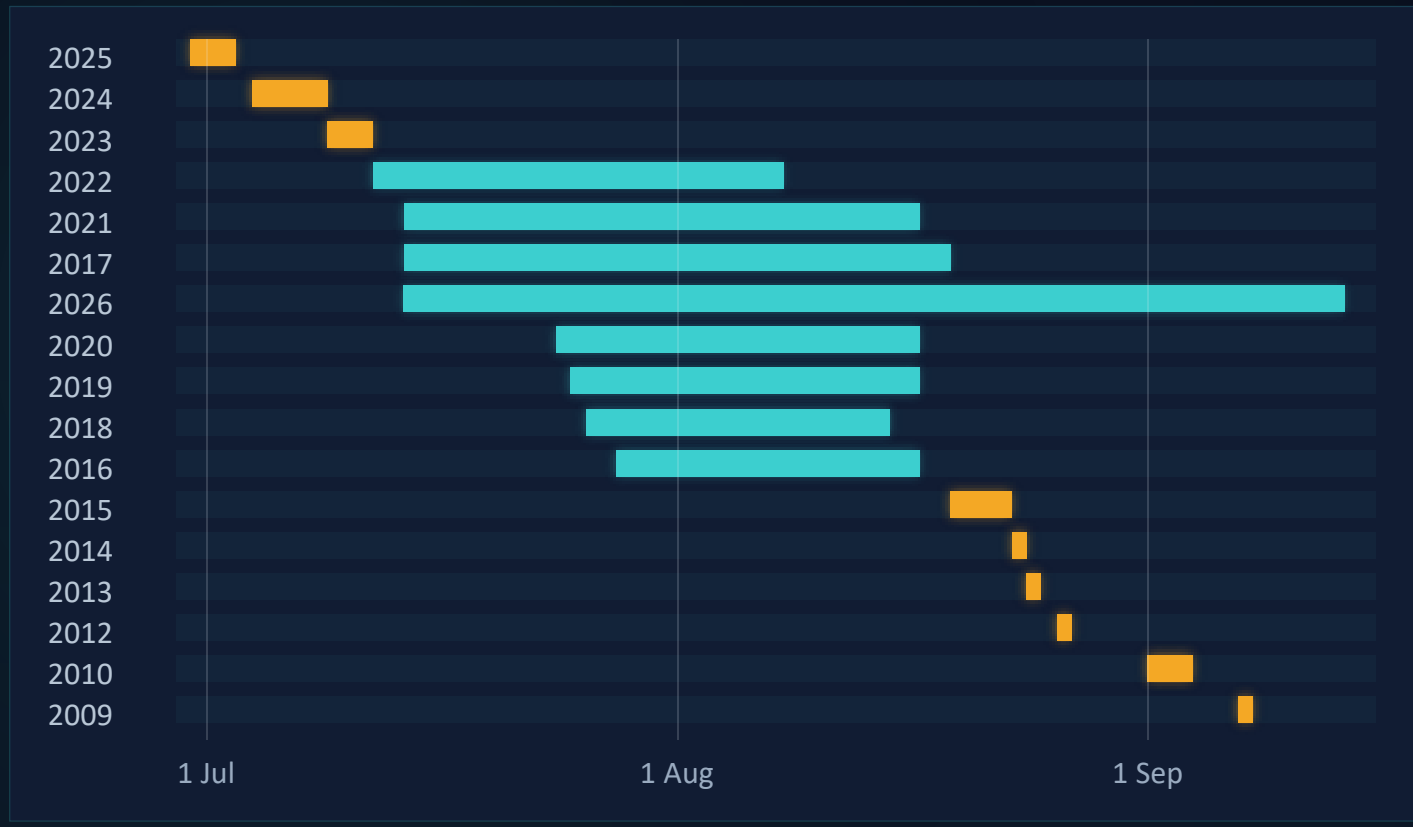

■ one production pass

■ several passes (gap recovery, fixes)

Rebuilding the network, year by year

STATIONS PER DAILY SOLUTION · MEDIAN, INCL. 14 TIES

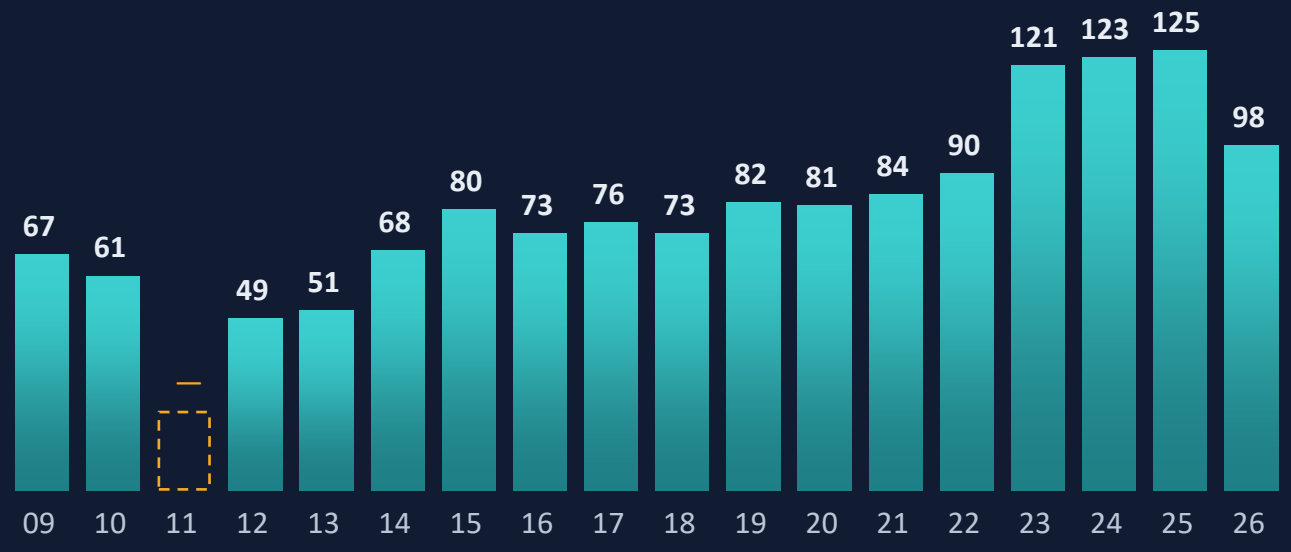

2009 → 2026 · the 2011 archive is lost

ONE FRAME FOR EVERY YEAR

Orbits & clocks
CODE repro3 to 2022 DOY 330, then operational

Antennas & frame
igs20.atx for all years · IGS20 datum

Daily datum
no-net-translation on 10–12 IGS stations

Recovered
hidden folders, 1 Hz gap filling, truncated files

One console for the whole operation

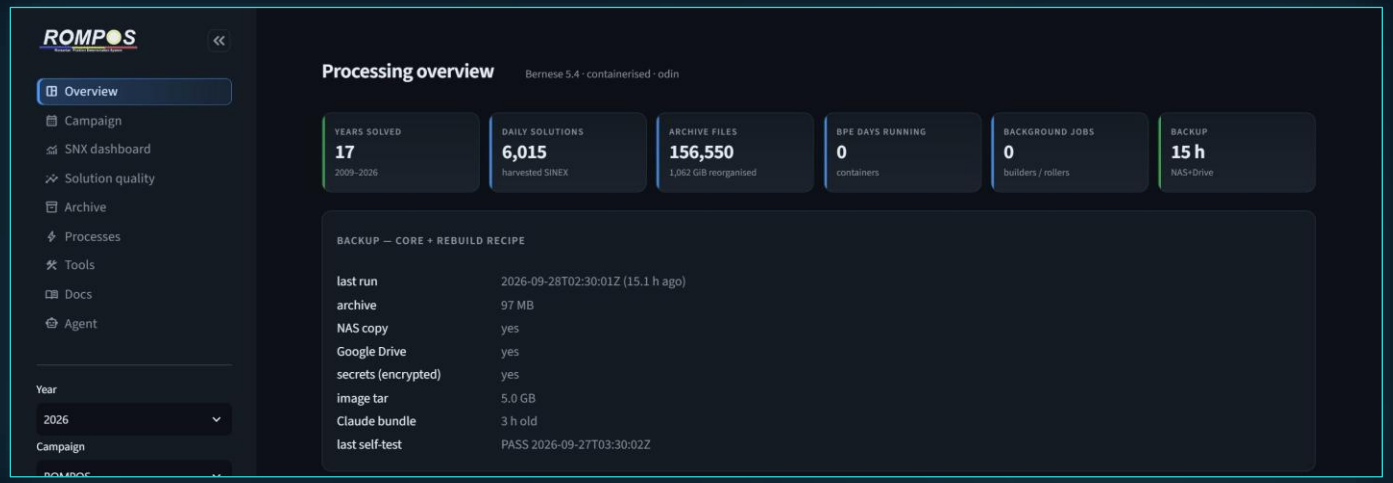

Streamlit 1.61
3,758 lines of Python

13 charts
6 matplotlib · 6 plotly · 1 native

3-layer cache
in memory + fingerprinted JSON on disk

NINE PAGES

Overview	KPIs, storage, backup health
Campaign	day grid for every year
SNX dashboard	8 analysis tabs
Solution quality	repeatability, variance factor
Archive	index and relay build progress
Processes	live containers, BPE step %
Tools	reprocess, deliver, STA from sitelogs
Docs	browse the knowledge base
Agent	Claude chat with the tools

2025 · 365 days · 132 stations

Every station, every day — at a glance

Punch card hours missing per station × day

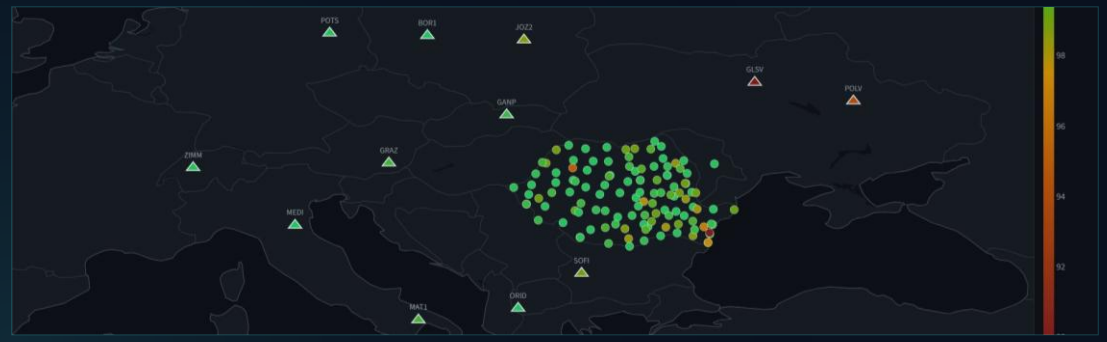

Station map coverage per site, IGS/EPN ties as triangles

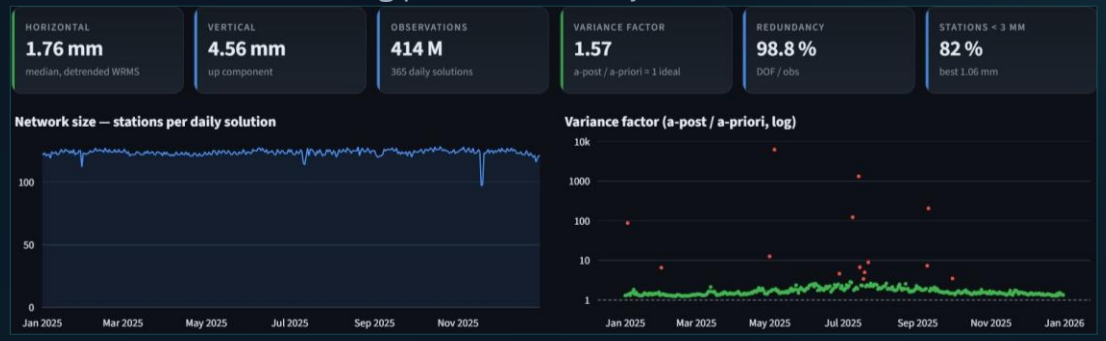

Solution quality WRMS, variance factor, network size

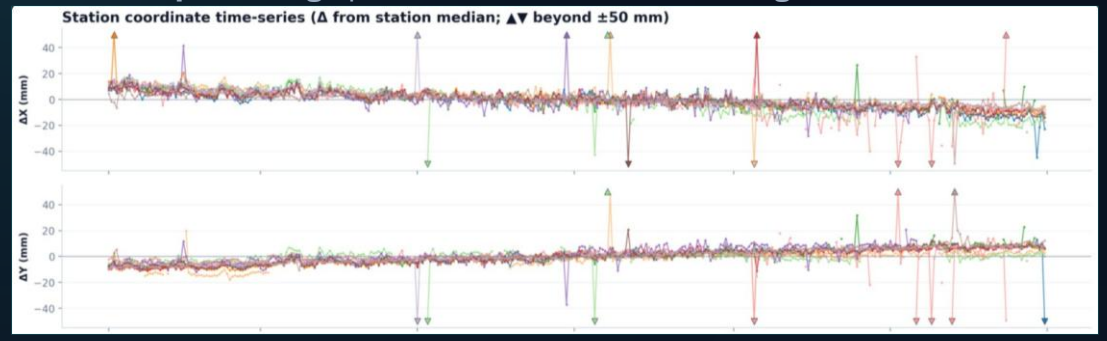

Coordinate series Δ from station median, outliers flagged

// PART 03

The AI layer

Agent, knowledge, tools — and who checks them

Ask the pipeline — don’t operate it

Operator
asks in plain language

↓

Claude agent
≤ 8 tool rounds · web search

↓

MCP server — 8 tools
status · query · quality · run

↓

Guardrail
writes need approval · audited

LIVE ANSWER FROM THE CONSOLE — GROUNDED IN THE KNOWLEDGE BASE

Why do the CODE repro3 orbit files say PCV:IGS14 in their header if we process in IGS20?

Short answer: the header string is correct — it's an **ATX lineage label, not a frame label.**

CODE repro3 orbits were computed with `igsR3.atx` **(a.k.a. R20.ATX)** — a distinct antenna model that is *neither* igs14.atx nor igs20.atx. Crucially, **repro3 predates IGS20**: the repro3 campaign finished in 2020–2021, ITRF2020 came out April 2022, and IGS20 only went operational at GPS week 2238 (27 Nov 2022). A 2021 campaign simply couldn't be expressed in a 2022 frame.

The `PCV:IGS14` string appears because `igsR3.atx` **internally declares** `ANTENNA MODEL: IGS14` — it belongs to the IGS14 lineage — and the SP3 `PCV:` field just echoes the ATX's own model string. We confirmed this on odin against our own Bernese conversion in `DATAPOOL/REF54/ANTENNA_R20.PCV`, which shows `ANTENNA MODEL: IGS14` / `LAST INPUT FILE: R20.ATX` (built 10-DEC-2020, SINEX tag `IGSR3_2135`). So the files are neither "wrong" nor in the old igs14.atx frame.

The chain is documented in IGSMAIL-8238 (26 Jul 2022): igs14.atx → **igsR3.atx** (major update, 37 legacy L1/L2 patterns replaced by multi-frequency Geo++/Bonn patterns) → **igs20.atx** (minor update, 15 more Geo++ patterns + ITRF2020-scale satellite z-PCOs). The `ITRF2020_to_IGS20.txt` offsets are explicitly the corrections "from igsR3.atx to igs20.atx." Per Joaquín (IGS/ACC), the value of repro3 is in the antennas/models used, not the underlying frame — and ITRF2014/ITRF2020 are close enough that ITRF2014 is a reasonable first approximation to ITRF2020.

claude-opus-4-8
up to 8 tool rounds per question

Grounded
RAG knowledge base + live web search

Writes gated
dry-run until a human approves

// 20 KNOWLEDGE (RAG)

What the agent has read

IGS Mail 2008 – 2026

3,052 messages — frame switches, product and server changes

Bernese announcements

427-entry AIUB index + full BSW messages

EPN / EUREF guidelines guidelines

analysis centres, stations, combination, ETRS89

repro3 science

IGS & EPN repro3 papers, GNSS modelling reviews

EPN-CB news

station and network news, refreshed weekly

Our own playbooks

architecture, decisions, every failure + fix

41 **documents**
7,248 chunks

- BM25 retrieval over headed chunks
- sources re-fetched every Monday
- live web search for anything newer
- answers cite their source

EXAMPLE

“Why do repro3 orbit files say PCV:IGS14?” → igsR3.atx declares ANTENNA MODEL: IGS14 — a lineage label, not the frame (IGSMAIL-8238).

// 21 MCP SERVER

Eight tools — the agent's hands

READ · free to use

status
processing state per year

query
read-only SQL on the state database

assess_quality
Helmert fit of a day vs IGS20

knowledge_search
search the RAG library

audit_log
who ran what, when

WRITE · gated

run_doy
process one day in a container

decimate
1 Hz → 30 s gap recovery

refresh_decimation
rolling 14-day refresh

every write: dry-run unless confirm=true + human approval

WHY MCP

One standard interface
Model Context Protocol: the same tool server serves the console agent and Claude Code.

One choke point
every write passes the same gate into audit.log — time, tool, arguments, executed or not.

Tools describe themselves
the agent reads the schema and writes its own SQL.

// 22 HUMAN + AI

Two agents, one operator

Operations agent
in the console

12
tool calls · all reads

- answers status, quality and ‘why’ questions
- writes SQL over the state database
- runs a day or a recovery — after approval
- cites the knowledge base and the web

Engineering agent
Claude Code

3,387
tool calls · 2,347 on odin

- built and fixed the whole pipeline over ssh
- runs campaigns and watches them
- multi-agent reviews: finders + adversarial verifiers
- persistent memory; asks before metadata changes

THE WHOLE CAMPAIGN WAS DRIVEN LIKE THIS — REAL COMMANDS

> process 2014

> lets prepare 2010

> for CRAI until 17.12 it was Ashtech

> stats?

> make a button harvest to deliver

> do the proposed fixes and reprocess affected days

Find → verify → fix → reprocess

Find
outliers in the coordinate series

↓

Verify
adversarial agents refute it

↓

Fix
metadata, links, coordinates — approved

↓

Reprocess
affected days, then QC again

12 agents diagnose · 8 more review the fix

CASE STUDY · 2026

68 mm	DYNG antenna replaced on 21 May, missing from station info → height bias
11 cm	a wrongly fixed link baseline shifted 10 stations (DOY 095; 9 on DOY 118)
1.25 m	wrong ambiguities on a 1-hour session (HUED, DOY 134)
0.8 m	wrong a priori coordinates: CALR 0.80 m, PTRS 0.45 m

49 days reprocessed · final-vs-float flags 30 → 12 · others moved ≤ 2 mm

Surviving Claude Code transcripts, 4 Jul – 29 Sep 2026 · lower bounds

The engineering agent, by the numbers

3,387
agent tool calls
from 381 operator prompts (median 6 words): ~9 actions each

4,810
shell commands on odin
2,347 by the main agent, 2,463 by its sub-agents

155
scripts on odin · 17,472 lines
91 of the 97 new since 4 Jul created by the AI (~11,300 lines)

9,907
Bernese day-runs
→ 6,015 daily solutions (99.85 %), 483,211 station-days

69,479
recovered station-days
from 1 Hz data, 2022–25; plus 103,593 legacy station-days rebuilt

9 · 74
workflows · sub-agents
2,702 sub-agent tool calls, 0 agent errors

15
decisions handed back
era split, marker identities, weekly-solution definition, …

~53 h
engaged agent time, 47 days
1.90 B tokens (95 % cache reads) · 4 Claude models

// 25 TIME-CRITICAL INCIDENTS

UTC, from the moment the problem was raised · raised by = who noticed first

Problems solved against the clock

TIME	DATE	INCIDENT → FIX	RESULT	RAISED BY
53 min	7 Jul	2024 ‘dark zone’: data declared lost sat under RINEX 3 short names pipeline rewired in 6 min, archive rescanned with every naming variant	+1,358 station-days	operator
19 min	11 Jul	odin off the network for 16 h: all jobs killed, NFS mounts gone after the power-cycle: remount and resume, from its own restart playbook	campaign resumed	AI
2 h 04	13 Jul	2022 roll: cleanup deleted a GPS week of orbits still in use stopped in 2.4 min, root cause in 37, relaunch after a 2-week test	2022: 365 / 365	operator
50 min	10 Aug	rc = 0 hid 522 empty BPE runs; 6 markers never in a solution root cause in 50 min; exit-code fix, PPP a priori, 2016–21 re-run	+7,293 station-days	AI audit
2 min	19 Aug	21 days of 2026 failing in the final GPSEST, called ‘unfixable’ for 35 days 16,072 ambiguities > 16,000 limit → cluster at 80; all 21 solved in 4 h	2026: 159 → 180 days	operator
1 h 55	19 Aug	2015 launch blocked by 4 silent pre-2016 defects dead AIUB FTP → CDDIS · staging gate · missing GIM · truncated .gz salvaged	2015: 365 / 365	AI
2.4 min	15 Sep	2 of 756 weekly SINEX truncated while ADDNEQ2 reported OK rebuilt, %ENDSNX gate added, 9 files re-uploaded	fixed before sending	AI

// 26 WHO CHECKS THE AI

Verified by agents, overruled by people

1 in 3

first-pass AI conclusions changed by adversarial verifiers

11 of 34 verdicts in 6 workflows refuted a finding or judged a fix unsound

HONEST LIMITS

- 3 AI misdiagnoses caught by the operator: ‘dark zone’, ‘unfixable’ 2026 days, COST ‘missing’
- ≥ 24 messages where the AI owned a mistake
- console agent: 12 tool calls, 0 writes — the work ran through Claude Code

GFZ DETOUR

51 min

The AI moved 2022 to GFZ repro3. Five skeptics found zero satellite biases (CODE: ~13,000). Reverted before a single day ran.

FIX REVIEW

4 of 4

fix recipes judged unsound — one hook would have failed every day’s processing. The corrected recipes were deployed.

HAR1

3 winters

A proposed +51 mm station break was refuted: the same height rise returns every winter. No break applied.

PETO ≠ PET2

351 m

The operator overruled a merge the AI had endorsed. PPP later put the two 351 m apart — 1,213 station-days saved.

// 27 RUNNING IT

Built to be operated — and rebuilt

Live processes

- every container and its BPE step %
- auto-refresh every 10 s
- logs one click away

10 s

refresh while jobs run

Harvest to deliver

- pre-analyses the years on the NAS
- missing GPS weeks built first
- SHA-256 for every file

1

button per delivery

3-2-1 backup

- NAS + Google Drive: 24 / 24 nights
- GPG-encrypted secrets, rebuild recipe
- Sunday self-tests: 4 / 4 PASS

02:30

every night · self-test weekly

// 28 LESSONS

What we learned

Measure before you tune
the limit was one single-threaded GPSEST, not the core count.

Verify — never trust a green light
Bernese said OK on a truncated SINEX; a 68 mm antenna bias raised only a warning.

Keep the human on the metadata
every station, coordinate and exclusion change: AI proposes, a person approves.

Write the knowledge down
RAG, playbooks, memory: each session starts where the last one ended.

NEXT

- balanced clusters for every day
- tune BPE jobs per day (Maxj) and stage in parallel
- 2008 — the last viable year
- 2026 processed live, day by day
- EUREF-style weekly submissions

// END OF TRANSMISSION

From a program you operate to a system you talk to.

Miluță DULEA-FLUERAȘ
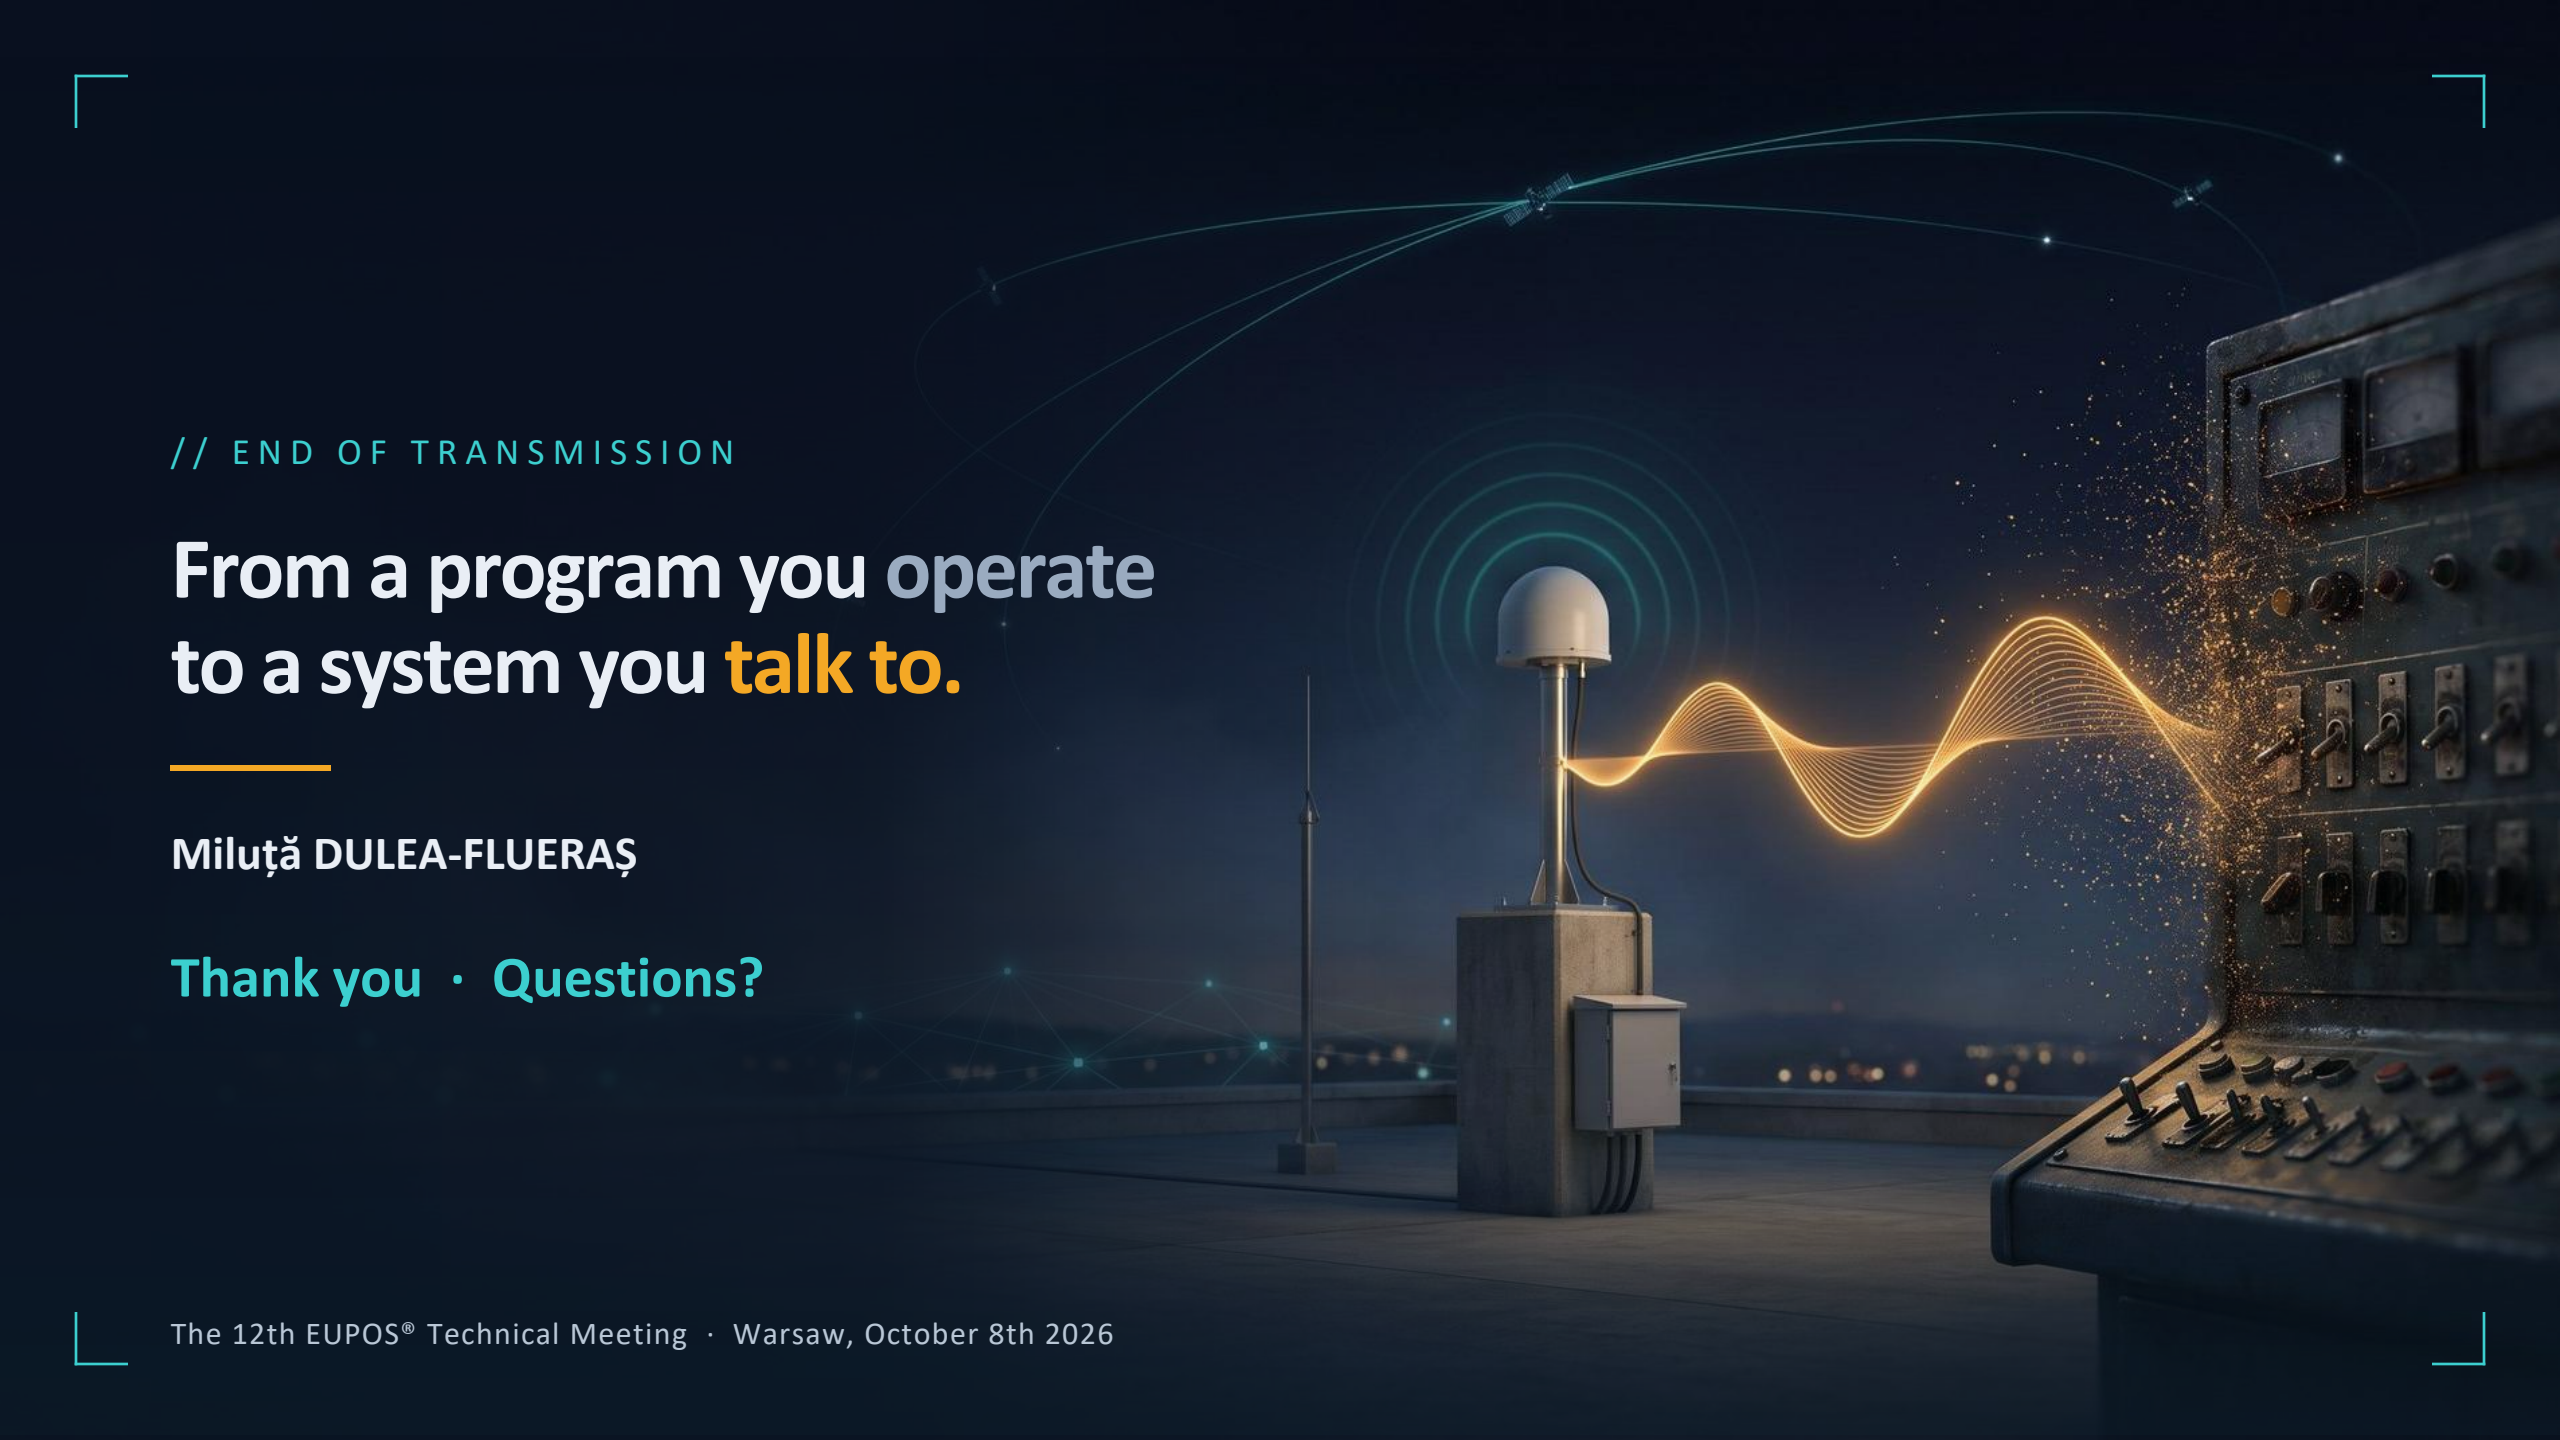
Thank you · Questions?